



Ubertas Consulting Limited
Cloud Hosting Service Definition Document

Table of Contents

2.0	AWS Security Assurance	3
2.1	Information Assurance	3
2.1.1	ISO 27001 Certification.....	3
2.1.2	NCSC UK Cloud Security Principles.....	4
2.2	GDPR and processing of Personal Data.....	4
2.3	AWS re:Start Program	4
3.0	Service Definition – AWS Config	6
3.1	Service Overview	6
3.2	Backup/Restore and Disaster Recovery	6
3.3	Pricing Overview	6
3.4	Service Constraints	6
3.5	Technical Requirements	7
4.0	Cross-Service Definitions	8
4.1	Availability	8
4.1.1	Region Availability	9
4.2	On-Boarding/Off-Boarding Processes and Service Migration	10
4.3	Service Management Details	10
4.4	Service Levels and Service Credits.....	11
4.5	Trial Service Details	15
4.6	Data Restoration/Service Migration.....	15
4.6.2	Service Migration.....	19
4.7	Customer Responsibilities.....	20



2.0 AWS Security Assurance

Moving IT infrastructure to AWS means that both the customer and AWS have important roles in the operation and management of security in their areas of responsibility. AWS operates, manages, and controls the components from the host operating system and virtualisation layer down to the physical security of the facilities in which the services operate. The customer is responsible for management of the guest operating system (including updates and security patches to the guest operating system) and associated application software, as well as the configuration of the AWS provided security group firewall and other security-related features. The customer will generally connect to the AWS environment through services provided by third parties (e.g., internet service providers). AWS does not provide these connections, and the customer should consider the security of such connections and the security responsibilities of such third parties in relation to their systems. This is really no different from working with a network service provider who brings connectivity to on-premises data centres.

We are vigilant about the security of our underlying cloud environment and have implemented sophisticated technical and organisational measures against unauthorised access. Customers can validate the security controls in place within the AWS environment through AWS certifications and reports, including the AWS System and Organisation Controls (SOC) 1, 2, and 3 reports, International Organisation for Standardization (ISO) 27001 certification, and Payment Card Industry Data Security Standard (PCI DSS) compliance reports. These reports and certifications are produced by independent third-party auditors and attest to the design and operating effectiveness of AWS security controls. The applicable AWS compliance certifications and reports can be requested at <https://aws.amazon.com/compliance/contact>. More information on AWS compliance certifications, reports, and alignment with best practices and standards can be found at <https://aws.amazon.com/compliance> and <https://aws.amazon.com/compliance/programs/>.

2.1 Information Assurance

The following subsections provide information relating to information assurance.

2.1.1 ISO 27001 Certification

AWS is certified under the ISO 27001 standard. ISO 27001 is a widely adopted global security standard that outlines the requirements for information security management systems. It provides a systematic approach to managing company and customer information that is based on periodic risk assessments. In order to achieve the certification, a company must show it has a systematic and ongoing approach to managing information security risks that affect the confidentiality, integrity, and availability of company and customer information.

AWS has established a formal programme to maintain the certification. More information regarding AWS's ISO 27001 certification can be found at <http://aws.amazon.com/compliance/iso-27001-faqs/>.

2.1.2 NCSC UK Cloud Security Principles

In 2016, National Cyber Security Centre (NCSC) UK published the [Cloud Security Collection](#) documents for public sector organisations that are considering the use of cloud services for handling information classified as OFFICIAL. The collection of guidance documents aims to help public sector organisations make informed decisions about cloud services and choose a cloud service that balances business benefits and security risks. In order to provide you with more information regarding NCSC UK's Cloud Security Principles and to make an informed decision when performing risk assessments, we have published a whitepaper called [Using AWS in the Context of NCSC UK's Cloud Security Principles](#).

This whitepaper provides insights into implementation and assurance approaches within AWS based on the published guidance for each of the 14 [Cloud Security Principles](#) and provides an in-depth view into the AWS implementation approach in relation to the Cloud Security Principles. Based on this information, UK public sector organisations and their information security functions can conduct informed risk assessments and select the appropriate AWS Cloud services for their cloud environment.

2.2 GDPR and processing of Personal Data

AWS offers a GDPR-compliant Data Processing Addendum (DPA), enabling customers to comply with GDPR contractual obligations. More information can be found at the following links:

- AWS GDPR Center: <https://aws.amazon.com/de/compliance/gdpr-center/>
- AWS EU Data Protection website: <https://aws.amazon.com/compliance/eudata-protection/>

2.3 AWS re:Start Program

AWS re:Start is an initiative intended to promote the development of much-needed technical skills in the UK. It is a training and job placement programme that was established in late March 2017, coincident with the launch of the AWS London Region. AWS re:Start has been designed to educate young adults, military veterans, members of the military reserve, those leaving the Armed Forces, and service spouses on the latest software development and cloud computing technologies.

AWS re:Start is the result of a close collaboration between the Ministry of Defence, QA Consulting (an APN Training Partner) and The Prince's Trust. In conjunction with members of the AWS Partner Network (APN) and customers, the programme aims to offer work placements to 1,000 people.

AWS re:Start accommodates participants at all levels of experience—even those with no previous technical knowledge. Participants who join AWS re:Start complete technical training classes, led by AWS certified instructors, and gain experience through on-the-job training. They learn about multi-tier architectures, application programming interfaces (APIs), and microservices. Training content for the AWS re:Start program is



curated by AWS in collaboration with QA Consulting, which delivers the training courses.

Organisations that have pledged job placements to AWS re:Start include Annalect, ARM, Claranet, Cloudreach, Direct Line Group, EDF Energy, Funding Circle, KCOM, Sage, Tesco Bank, and Zopa. Participants completing the programme are eligible for many different technical positions within these companies, including sought-after entry level positions such as first-line help desk support, IT support analyst, software developer, IT support technician, network engineer, IT recruitment consultant, and IT sales roles.

The programme also provides participants with the fundamental knowledge needed to start working with AWS to build technology start-up businesses. The Young Adults thread of AWS re:Start is delivered via the Prince's Trust Get into Technology initiative. In addition to technical training, the "Get into Technology" programme supports students with mentoring, soft work skills, and help in applying for jobs, including resume writing and interview skills. [Contact us here](#) for information and/or to get involved in the program.

3.0 Service Definition – AWS Config

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

3.1 Service Overview

AWS Config is a fully managed service that provides you with an AWS resource inventory, configuration history, and configuration change notifications to enable security and governance. With AWS Config you can discover existing AWS resources, export a complete inventory of your AWS resources with all configuration details, and determine how a resource was configured at any point in time. These capabilities enable compliance auditing, security analysis, resource change tracking, and troubleshooting.

Top benefits include:

- **Configuration Visibility** – You can view continuously updated details of all configuration attributes associated with AWS resources. You are notified via Amazon SNS of every configuration change and you can process these notifications programmatically.
- **Fully Managed** – With AWS Config, there are no software agents to install, and no databases to manage. AWS Config automatically manages the underlying infrastructure required to record, store and report on the configuration details of your AWS resources.
- **Easy to get started** – You can enable AWS Config with a few clicks in the AWS Management Console. AWS Config will discover your AWS resources and start recording configuration changes. You can access information about the configuration of any resource using the AWS Management Console, CLI, or SDKs.
- **Low cost** – With AWS Config, there are no upfront costs. You pay-as-you-go based on the number of resources and configuration changes recorded for your AWS account.

3.2 Backup/Restore and Disaster Recovery

AWS Config uses an Amazon S3 bucket you specify to store the information recorded. Refer to Amazon S3 for more detailed information.

3.3 Pricing Overview

Please visit <https://aws.amazon.com/pricing/> for further details on pricing AWS services.

3.4 Service Constraints

Please see <https://aws.amazon.com/config/> for more information.

3.5 Technical Requirements

AWS Config provides a detailed view of the resources associated with your AWS account, including how they are configured, how they are related to one another, and how the configurations, and their relationships have changed over time.

- **Developer Guide** – Provides a conceptual overview of AWS Config and includes detailed development instructions for using the various features. [HTML](#) | [PDF](#)
- **CLI Reference** – Documents the AWS Config CLI. [HTML](#)
- **API Reference** – Describes all the API operations for AWS Config in detail. Also provides sample requests, responses, and errors for the supported web services protocols. [HTML](#) | [PDF](#)

4.0 Cross-Service Definitions

The following service definition topics are applicable to all AWS Service Offerings and are detailed once in a cross-service manner below.

4.1 Availability

AWS Cloud services are hosted within our global data centre footprint, allowing customers to consume services without having to build or manage facilities or equipment. AWS Cloud services are offered in separate Regions in a number of separate geographic areas. A Region is a physical location in the world where we have multiple, isolated locations known as Availability Zones that are engineered to be isolated from failures in other zones (see **Figure 1** below). Availability Zones consist of one or more discrete data centres, each with redundant power, networking, and connectivity, and housed in separate facilities.




Figure 1 – Each Availability Zones can consist of multiple data centres, and at full scale can contain hundreds of thousands of servers. Every AWS Region contains 2+ zones and some Regions have as many as 6 zones

Availability Zones are engineered to be isolated from failures in other zones, and to provide inexpensive, low-latency network connectivity to other zones in the same Region. By hosting workloads in separate zones, you can protect your applications from the failure of a single location. Availability Zones offer the ability to operate production applications and databases that are more highly available, fault tolerant, and scalable than would be possible with a single data centre.

Customers have access to 22 AWS Regions around the globe, including five Regions in the EU—United Kingdom (London), Ireland (Dublin), Germany (Frankfurt), France (Paris), and Sweden (Stockholm)—with additional EU Regions (located in Milan, Italy and in Spain) coming soon. Customers can choose to use one Region, all Regions, or any combination of Regions. Together, our EU Regions allow customers to architect highly fault tolerant applications while storing their data in the EU.

The AWS Cloud provides customers with the flexibility to run workloads and store data within multiple geographic Regions as well as across multiple Availability Zones within each Region. You decide which AWS Region(s) house your data, and it resides only in

the Region(s) you specify, for as long as you choose. For example, a customer can choose to deploy their AWS Cloud services and data exclusively in the London Region, and customer content is not moved outside of London unless the customer decides to move it.

We are steadily expanding global infrastructure to help our customers achieve lower latency and higher throughput. As our customers grow their businesses, AWS will continue to provide infrastructure that meets their global requirements. The AWS products and services that are available in each Region are listed in [our Region Table](#).

Figure 2 displays our 22 global Regions and 69 Availability Zones (and one local Region in Osaka, Japan¹). Five more AWS Regions in Cape Town, Jakarta, Milan, Osaka, and Spain have been announced.

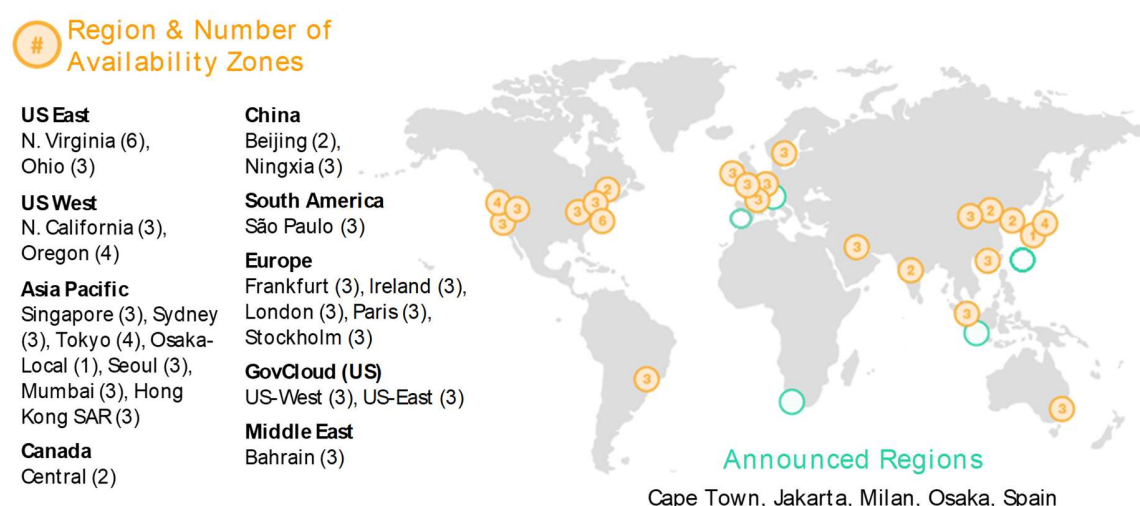


Figure 2 – AWS's Global Infrastructure Consists of 22 Regions and 69 Availability Zones

To deliver content to end users with lower latency, Amazon CloudFront uses a global network of 187 Points of Presence (176 Edge Locations and 11 Regional Edge Caches) in 69 cities across 30 countries. Visit our website for a list of current [Amazon CloudFront edge locations](#).

4.1.1 Region Availability

Exact service availability depends on a range of factors and choices made by customers when they architect and implement their solution.

The Services Offerings will be delivered from the AWS Region selected by the customer upon opening an AWS account. The customer may specify the AWS Region in which customer content will be stored. It is the customer's responsibility to select the relevant

¹ For customers who specifically need to replicate their data or applications over greater geographic distances, there are AWS Local Regions. An AWS Local Region is a single datacentre designed to complement an existing AWS Region. Like all AWS Regions, AWS Local Regions are completely isolated from other AWS Regions.



AWS Region in order to comply with its own security and governance requirements. AWS will not access or use customer content except as necessary to maintain or provide the Service Offerings, or as required by law or regulation. Customers acknowledge that AWS does not limit customers to any particular AWS Region. Note that not all AWS Cloud services are available in every AWS Region; however, we are steadily expanding our service availability across AWS's global regions.

The full list of available AWS services, and their availability by region can be seen on our website at <https://aws.amazon.com/about-aws/global-infrastructure/regional-product-services/>.

4.2 On-Boarding/Off-Boarding Processes and Service Migration

AWS maintains a cadre of Getting Started Guides and schedules regular webinars. These guides and webinars cover a variety of topics, including the two discussed in the following bullets. See <http://aws.amazon.com/documentation/gettingstarted/> for more details.

- **Getting Started with AWS** – This guide provides an introduction to AWS, examples of what you can do with AWS, basic information that you need to know to get started, and links to resources and documentation that will help you learn more no matter what your use case is. The guide is available in [HTML](#), [PDF](#), and [Kindle](#) formats.
- **AWS Management Console** – This guide provides an overview of the AWS Management Console and instructions on starting various services, steps to customising the navigation bar, and tips on accessing region settings, security credentials, and billing updates. The guide is available in [HTML](#) format.

AWS allows customers to move data as needed off AWS storage using the public internet or AWS Cloud services such as AWS Direct Connect, AWS Import/Export, and more.

With AWS, you can provision compute power, storage, and other resources, gaining access to a suite of elastic IT infrastructure services as your business demands them. With minimal cost and effort, you can move your application to the AWS Cloud and reduce capital expenses, minimise support and administrative costs, and retain the performance, security, and reliability requirements your business demands. To see a step-by-step migration strategy, refer to the [Migrating Your Existing Applications to the AWS Cloud](#) whitepaper.

4.3 Service Management Details

All AWS Cloud services are driven by robust APIs that allow for a wide variety of monitoring and management tools to integrate easily with your AWS Cloud resources.

Common tools from vendors such as Microsoft, VMware, BMC Software, Okta, RightScale, Eucalyptus, CA, Xceedium, Symantec, Racemi, and Dell, to name just a few, already support AWS.

4.3.1.1 **AWS Management Console**

The AWS Management Console is a single destination for managing all of your AWS resources, from Amazon EC2 instances to Amazon DynamoDB tables. Use the console to perform any number of tasks, from deploying new applications to monitoring the health of your applications. The console enables you to manage all aspects of your AWS account, including accessing your monthly spending by service, managing security credentials, or even setting up new IAM users. The console supports all AWS Regions and lets you provision resources across multiple regions.

4.3.1.2 **AWS Command Line Interface (AWS CLI)**

The AWS CLI is a unified tool used to manage your AWS Cloud services. With just one tool to download and configure, you can control multiple AWS Cloud services from the command line and automate them through scripts.

4.3.1.3 **Use Your Existing Management Tools**

It is likely that many of the tools that your organisation is using to manage your on premises environments can be integrated with AWS. Integrating your AWS environment can provide a simpler and quicker path for cloud adoption, because your operations team does not need to learn new tools or develop completely new processes.

4.4 **Service Levels and Service Credits**

AWS currently provides SLAs, with a corresponding Service Credit regime, for several products. Due to the rapidly evolving nature of AWS's product offerings, SLAs are best reviewed directly on our website via the links below:

[Alexa for Business Service Level Agreement](#)

[Amazon AppStream 2.0 Service Level Agreement](#)

[Amazon API Gateway Service Level Agreement](#)

[Amazon Athena Service Level Agreement](#)

[Amazon Aurora Service Level Agreement](#)

[Amazon Chime Service Level Agreement](#)

[Amazon Cloud Directory Service Level Agreement](#)

[Amazon CloudFront Service Level Agreement](#)

[Amazon CloudSearch Service Level Agreement](#)

[Amazon CloudWatch Service Level Agreement](#)

[Amazon Cognito Service Level Agreement](#)

[Amazon Compute Service Level Agreement](#)

[Amazon Connect Service Level Agreement](#)
[Amazon DocumentDB \(with MongoDB compatibility\) Service Level Agreement](#)
[Amazon DynamoDB Service Level Agreement](#)
[Amazon EC2 Service Level Agreement](#)
[Amazon EFS Service Level Agreement](#)
[Amazon EKS Service Level Agreement](#)
[Amazon Elastic Container Registry Service Level Agreement](#)
[Amazon Elastic Load Balancing Service Level Agreement](#)
[Amazon Elastic Transcoder Service Level Agreement](#)
[Amazon ElastiCache Service Level Agreement](#)
[Amazon Elasticsearch Service - Service Level Agreement](#)
[Amazon EMR Service Level Agreement](#)
[Amazon FSx Service Level Agreement](#)
[Amazon GuardDuty Service Level Agreement](#)
[Amazon Inspector Service Level Agreement](#)
[Amazon Kinesis Service Level Agreement](#)
[Amazon Lightsail Instance and Block Storage Service Level Agreement](#)
[Amazon Lightsail Managed Databases Service Level Agreement](#)
[Amazon Macie Service Level Agreement](#)
[Amazon Machine Learning Language Service Level Agreement](#)
[Amazon Managed Blockchain Service Level Agreement](#)
[Amazon Messaging \(SQS, SNS\) Service Level Agreement](#)
[Amazon MQ Service Level Agreement](#)
[Amazon Neptune Service Level Agreement](#)
[Amazon QuickSight Service Level Agreement](#)
[Amazon RDS Service Level Agreement](#)
[Amazon Redshift Service Level Agreement](#)
[Amazon Rekognition Service Level Agreement](#)
[Amazon Route 53 Service Level Agreement](#)
[Amazon S3 Service Level Agreement](#)
[Amazon SageMaker Service Level Agreement](#)
[Amazon Simple Workflow Service Level Agreement](#)

[Amazon SimpleDB Service Level Agreement](#)
[Amazon User Engagement \(Pinpoint, SES\) Service Level Agreement](#)
[Amazon VPC NAT Gateway Service Level Agreement](#)
[Amazon WorkDocs Service Level Agreement](#)
[Amazon WorkLink Service Level Agreement](#)
[Amazon WorkMail Service Level Agreement](#)
[Amazon WorkSpaces Service Level Agreement](#)
[AWS Amplify Console Service Level Agreement](#)
[AWS AppSync Service Level Agreement](#)
[AWS Backup Service Level Agreement](#)
[AWS Budgets Service Level Agreement](#)
[AWS Certificate Manager Private Certificate Authority Service Level Agreement](#)
[AWS Client VPN Service Level Agreement](#)
[AWS CloudHSM Service Level Agreement](#)
[AWS Cloud Map Service Level Agreement](#)
[AWS CloudTrail Service Level Agreement](#)
[AWS CodeBuild Service Level Agreement](#)
[AWS CodeCommit Service Level Agreement](#)
[AWS CodeDeploy Service Level Agreement](#)
[AWS CodePipeline Service Level Agreement](#)
[AWS Config Service Level Agreement](#)
[AWS Cost Explorer API Service Level Agreement](#)
[AWS Database Migration Service Level Agreement](#)
[AWS Data Pipeline Service Level Agreement](#)
[AWS Device Farm Service Level Agreement](#)
[AWS Direct Connect Service Level Agreement](#)
[AWS Directory Service Level Agreement](#)
[AWS Elemental MediaConnect Service Level Agreement](#)
[AWS Elemental MediaConvert Service Level Agreement](#)
[AWS Elemental MediaLive Service Level Agreement](#)
[AWS Elemental MediaPackage Service Level Agreement](#)
[AWS Elemental MediaStore Service Level Agreement](#)

[AWS Elemental MediaTailor Service Level Agreement](#)
[AWS Firewall Manager Service Level Agreement](#)
[AWS GameLift Service Level Agreement](#)
[AWS Global Accelerator Service Level Agreement](#)
[AWS Glue Service Level Agreement](#)
[AWS Hybrid Storage and Data Transfer Service Level Agreement](#)
[AWS IoT 1-Click Service Level Agreement](#)
[AWS IoT Analytics Service Level Agreement](#)
[AWS IoT Core Service Level Agreement](#)
[AWS IoT Device Defender Service Level Agreement](#)
[AWS IoT Device Management Service Level Agreement](#)
[AWS IoT Greengrass Service Level Agreement](#)
[AWS Key Management Service Service Level Agreement](#)
[AWS Lambda Service Level Agreement](#)
[AWS OpsWorks Service Level Agreement](#)
[AWS PrivateLink Service Level Agreement](#)
[AWS RoboMaker Service Level Agreement](#)
[AWS Secrets Manager Service Level Agreement](#)
[AWS Security Hub Service Level Agreement](#)
[AWS Service Catalog Service Level Agreement](#)
[AWS Shield Advanced Service Level Agreement](#)
[AWS Site-to-Site VPN Service Level Agreement](#)
[AWS Step Functions Service Level Agreement](#)
[AWS Systems Manager Service Level Agreement](#)
[AWS Transit Gateway Service Level Agreement](#)
[AWS WAF Service Level Agreement](#)
[AWS X-Ray Service Level Agreement](#)

See the Supplier Terms document affiliated with this framework catalogue for additional information.

4.5 Trial Service Details

The AWS Free Tier is designed to enable you to get hands-on experience with AWS at no charge for 12 months after you sign up. After creating your AWS account, you can use products and services listed at <http://aws.amazon.com/free/> for free within certain usage limits.

4.6 Data Restoration/Service Migration

Traditional enterprise backup and recovery strategies typically take an agent-based approach whereby the entire contents of a server are backed up over either the Local Area Network (LAN) or the Storage Area Network (SAN). Conventional architectures have required this approach because replacing failed components is complex, time consuming, and operationally intensive. This has, in turn, created a backup environment that is complex to manage and resource intensive to operate, requiring technologies such as data de-duplication and virtual tape libraries to cope with ever-increasing workloads.

The AWS Cloud enables a far more lightweight approach to backup and recovery due, in part, to the following characteristics:

- Computers are now virtual abstract resources instantiated via code rather than being hardware based.
- Capacity is available at incremental cost rather than upfront cost.
- Resource provisioning takes place in minutes, lending itself to real-time configuration.
- Server images are available on demand, can be maintained by an organisation, and can be activated immediately.

These characteristics offer you opportunities to recover deleted or corrupted data, with less infrastructure overhead.

4.6.1.1 *Protecting Configurations Rather Than Servers*

Amazon Elastic Compute Cloud (Amazon EC2) enables the backup and recovery of a standard server, such as a web server or application server, so that you can focus on protecting configuration and stateful data rather than on the server itself. This set of data is much smaller than the aggregate set of server data, which typically includes various application files, operating system files, temporary files, and so on. This change of approach means that regular nightly incremental or weekly full backups can take far less time and consume less storage space.

When a compute instance is started in Amazon EC2, it is based upon an AMI and can also connect to existing storage volumes (e.g., Amazon EBS). In addition, when launching a new instance, it is possible to pass user data to the instance so that it can be accessed internally as dynamic configuration parameters.

A sample workflow would include the following steps:

1. Launch a new instance of a web server, passing it the “identity” of the web server and any security credentials required for initial setup. The instance is based upon a prebuilt AMI that contains the operating system and relevant web server application (e.g., Apache, Internet Information Services [IIS]).
2. Upon start-up, a boot script accesses a designated and secured Amazon S3 bucket that contains the specified configuration file(s).
3. The configuration file(s) contains various instructions for setting up the server (e.g., web server parameters, locations of related servers, additional software to install, patch updates).
4. The server executes the specified configuration and is ready for service. An open-source tool for performing this process, called cloud-init, is already installed on Amazon Linux AMIs and is also available for a number of other Linux distributions.

The figures below depict a traditional backup approach, and depicts an Amazon EC2 backup approach.

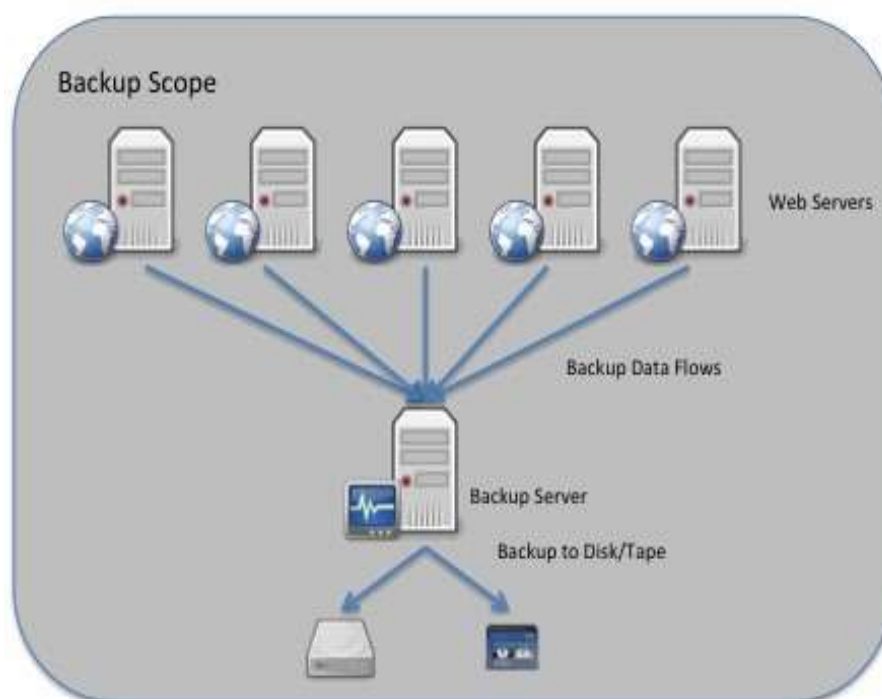


Figure 3 - Traditional Backup Approach

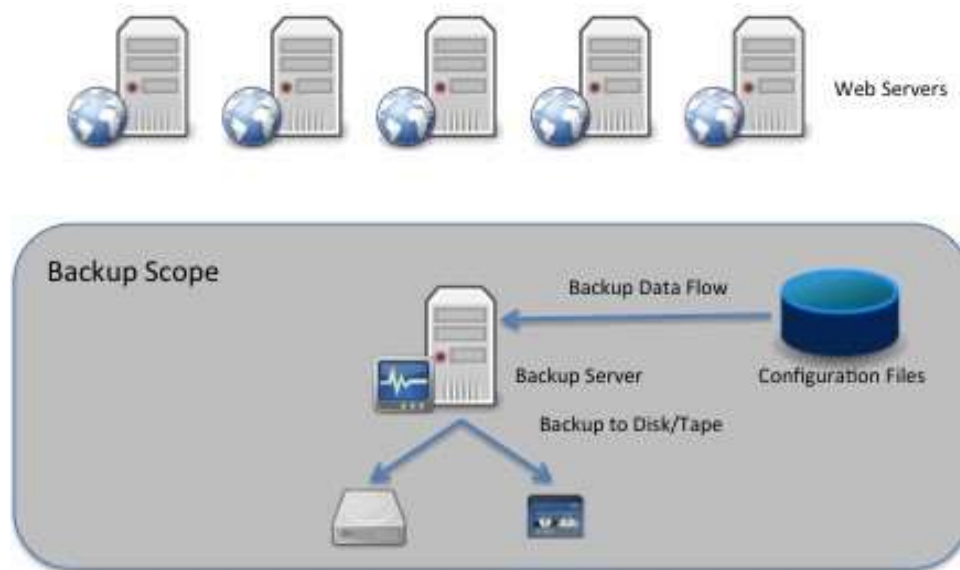


Figure 4 - Amazon EC2 Backup Approach

In the Amazon EC2 backup approach, there is no need to back up the server itself. The relevant configuration is contained in the combination of the AMI and the configuration file(s). Therefore, the only components requiring backup and recovery are the AMI and configuration file(s).

4.6.1.2 **Amazon Machine Images (AMIs)**

AMIs that you register are automatically stored in your account using Amazon EBS snapshots. These snapshots reside in Amazon S3 and are highly durable. This means that the underlying storage mechanism for the AMIs is protected from multiple failure scenarios.

It is also possible to share AMIs between separate AWS accounts. Consequently, you can create totally independent copies of the AMI by:

- Sharing the original AMI to another specified AWS account that you control.
- Starting a new instance based upon the shared AMI.
- Creating a new AMI from that running instance.
- The new AMI is then stored in the second account and is an independent copy of the original AMI. You also have the option to create multiple copies of the AMI within the same account.

4.6.1.3 **Configuration Files**

Customers use a variety of version management approaches for configuration files, and you can follow the same regime for the files used to configure your Amazon EC2 instances. For example, you could store different versions of configuration files in designated locations and securely control them like any other code. You could then

back up these code repositories using the appropriate backup cycle (e.g., daily, weekly, monthly) and snapshots to protected locations. Furthermore, you could use Amazon S3 to store your configuration files, taking advantage of the durability of the service in addition to backing up the files to an alternate location on a regular basis.

4.6.1.4 **Database and File Servers**

Backing up data for database and file servers differs from the web and application layers. In general, database and file servers contain larger amounts of business data (tens of GB to multiple TB) that must be retained and protected at all times. In these cases, you can leverage efficient data movement techniques such as snapshots to create backups that are fast, reliable, and space efficient.

For databases that are built upon Redundant Array of Independent Disks (RAID) sets of Amazon EBS volumes (and have total storage less than 1 TB), an alternative backup approach is to asynchronously replicate data to another database instance built using a single Amazon EBS volume. While the destination Amazon EBS volume will have slower performance, it is not being used for data access and can be easily snapshotted to Amazon S3 using the Amazon EBS snapshot capability.

More information on AWS backup and recovery can be found at <http://aws.amazon.com/backup-storage/>. Refer to the [Storage Options on the AWS Cloud](#) whitepaper for additional information.

4.6.1.5 **Using AWS for Disaster Recovery**

AWS provides the features and services that you can leverage for your disaster recovery processes. Using these AWS features, customers (oftentimes assisted by AWS partners) can build and operate agile and cost-effective disaster recovery solutions.

Traditional disaster recovery approaches involve the duplication of infrastructure to ensure the availability of spare capacity in a disaster scenario. This infrastructure needs to be procured, installed, and maintained so that it is ready to deal with the anticipated capacity requirements. Under normal operational circumstances, this infrastructure would typically be underutilised or over-provisioned.

With AWS, you can scale up your infrastructure on an as-needed basis, enabling faster disaster recovery of critical IT systems without incurring the infrastructure expense of a second physical site. This allows greater agility to change and optimise resources during a disaster recovery scenario. It also results in significant cost savings, because you only pay for what you use when using the highly scalable, reliable, secure, fast, and inexpensive AWS Cloud infrastructure.

The AWS Cloud supports many popular disaster recovery architectures, from “pilot light” environments that are ready to scale up at a moment’s notice to “hot standby” environments that enable rapid failover. Benefits of using AWS for disaster recovery include:

- **Fast Performance** – Readily access fast, disk-based storage and file retrieval.

- **No Tape** – Eliminate costs associated with transporting, storing, and retrieving tape media and associated tape backup software.
- **Compliance** – Avoid fines for missing compliance deadlines with fast retrieval of files.
- **Elasticity** – Quickly add any amount of data and easily expire and delete data without handling physical media.
- **Secure** – Trust our secure and durable technology platform that maintains industry-recognised certifications and audits.
- **Partners** – Connect with AWS solution providers and system integration partners to help with your deployment.

A business typically decides on an acceptable Recovery Time Objective (RTO) and Recovery Point Objective (RPO) based on the financial impact to the business when systems are unavailable. AWS can work with customers to plan disaster recovery solutions so that they cost-effectively provide system recovery based on the RPO within the timeline and service level established by the RTO. Additional information on using AWS for disaster recovery can be found at <http://aws.amazon.com/disaster-recovery/>. We also recommend reviewing the [Using AWS for Disaster Recovery](#) whitepaper for additional information.

4.6.2 Service Migration

AWS migration services include:

- **AWS Application Discovery Service** – A service that helps you quickly and reliably plan application migration projects by automatically identifying applications running in on-premises data centres, their associated dependencies, and their performance profile.
- **AWS Database Migration Service (AWS DMS)** – A service that helps you migrate databases to AWS easily and securely. The source database remains fully operational during the migration, minimizing downtime to applications that rely on the database. AWS DMS can migrate your data to and from most widely used commercial and open-source databases.
- **AWS Server Migration Service (AWS SMS)** – An agentless service that makes it easier and faster for you to migrate thousands of on-premises workloads to AWS. AWS SMS allows you to automate, schedule, and track incremental replications of live server volumes, making it easier for you to coordinate largescale server migrations.
- **AWS Snowball** – A petabyte-scale data transport solution that uses secure appliances to transfer large amounts of data into and out of the AWS Cloud.

Snowball addresses common challenges with large-scale data transfers, including high network costs, long transfer times, and security concerns. Transferring data with Snowball is simple, fast, secure, and can be as little as one-fifth the cost of high-speed internet.

- **AWS Snowball Edge** – A 100 TB data transfer device with on-board storage and compute capabilities. You can use Snowball Edge to move large amounts of data into and out of AWS, as a temporary storage tier for large local datasets, or to support local workloads in remote or offline locations.

4.7 Customer Responsibilities

As cloud computing customers are building systems on top of cloud infrastructure, the security and compliance responsibilities are shared between the Cloud Service Provider (CSP) and cloud customers. In an Infrastructure as a Service (IaaS) model, customers control how they architect and secure their applications and data put on the infrastructure, while CSPs are responsible for providing services on a highly secure and controlled platform and providing a wide array of additional security features. The level of CSP and customer responsibilities in this shared responsibility model depends on the cloud deployment model (see the [NIST Definition of Cloud Computing](#) models). Customers should be clear as to their responsibilities in each model. AWS's shared responsibility model is depicted in **the Figure 5** below.

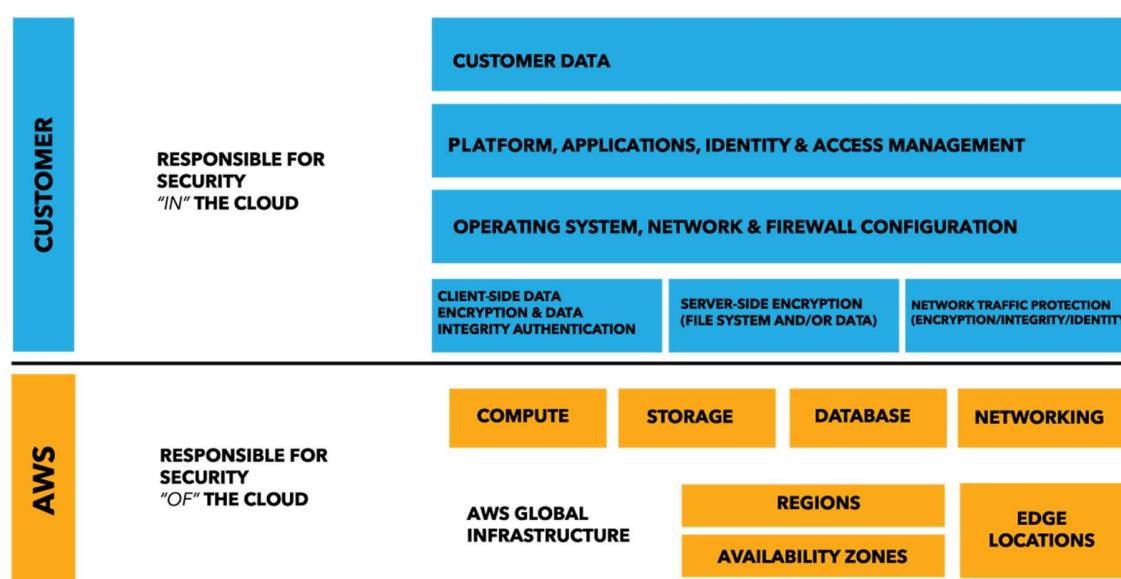


Figure 5 - AWS Shared Responsibility Model



AWS Responsibility – AWS operates, manages, and controls the infrastructure components, from the host operating system and virtualisation layer down to the physical security of the facilities in which the services operate.

Customer/Partner Responsibility – Customers/partners assume responsibility and management of the guest operating system (including updates and security patches), other associated application software, configuration of the AWS-provided security group firewalls, and other security, change management, and logging features.

AWS's shared responsibility model is further explained on the [AWS Compliance](#) webpage. AWS does not access customer data, and customers are given the choice as to how they store, manage, and protect their data. There are four important basics regarding data ownership and management in the shared responsibility model:

- Customers continue to own their data.
- Customers choose the geographic location(s) in which to store their data—it does not move unless the customer decides to move it.
- Customers can download or delete their data whenever they like.
- Customers should consider the sensitivity of their data and decide if and how to encrypt the data while it is in transit and at rest.